

(19)



Europäisches Patentamt

European Patent Office

Office européen des brevets



(11)

EP 0 892 335 A1

(12)

## EUROPEAN PATENT APPLICATION

(43) Date of publication:

20.01.1999 Bulletin 1999/03

(51) Int Cl.<sup>6</sup>: G06F 1/14, G06F 9/46

(21) Application number: 98304692.1

(22) Date of filing: 15.06.1998

(84) Designated Contracting States:

AT BE CH CY DE DK ES FI FR GB GR IE IT LI LU  
MC NL PT SE

Designated Extension States:

AL LT LV MK RO SI

(30) Priority: 07.07.1997 US 888551

(71) Applicant: INTERNATIONAL BUSINESS  
MACHINES CORPORATION  
Armonk, NY 10504 (US)

(72) Inventors:

- Berry, Robert F.  
Austin, Texas 78759 (US)
- Gu, Weiming  
Austin, Texas 78729 (US)

(74) Representative: Litherland, David Peter  
IBM United Kingdom Limited  
Intellectual Property Department  
Hursley Park  
Winchester, Hampshire SO21 2JN (GB)

(54) System and method for mapping processor clock values in a multiprocessor system

(57) A system and method is disclosed for accurately and efficiently synchronizing and then mapping, or normalizing, processor clocks in a multiprocessor information handling system. The system and method of the present invention provide sufficient granularity for sub-cycle variations between processors, while taking into account the problem of clock drifts. A plurality of processors are selected for the purpose of synchronization. The clocks located on the processors are synchronized, and then time values between synchronization points are mapped from each secondary processor to an equivalent, or normalized, time value in a primary processor. To accomplish this mapping, three clock differences are calculated. The first clock difference is the time between the first and second synchronization points for the primary processor, and the second clock difference is the time between the first and second synchronization points for the secondary processor. The third clock difference is the time difference between the time value to be mapped in the secondary processor and the first synchronization point in the secondary processor. The third time difference is multiplied by the ratio of the first time difference to the second time difference, and then the result is added to the time value of the primary clock at the first synchronization point. The final result is the mapped, or normalized, time value.




FIG. 1

## Description

The present invention relates to information handling systems and, more particularly, to a system and method for mapping, or normalizing, synchronized processor clocks in a multiprocessor information handling system.

In a multiprocessor information handling system, each processor may have its own independent clock, where each clock is driven by a different oscillator. Although the clocks are designed to run at a specific rate, each may run at a slightly different speed. In addition, external factors, such as room temperature, may affect the accuracy of the oscillators. Therefore, the clocks tend to "drift." In other words, the clocks are running at different and uneven speeds. Over a period of time, such drifts can become large enough to cause problems for applications that depend on consistent clock readings for all processors in the system.

There are many circumstances where processor clocks must be synchronized to perform certain system functions. One example is the sharing of performance information among the multiple processors in the information handling system. Typically, there is no means provided at the hardware level of an information handling system to provide for clock synchronization. Therefore, the operating system must provide for synchronization for its own processes. In addition to synchronizing the various processor clocks in the system, any clock readings taken must be "normalized," or adjusted, to account for the clock drift problem discussed above.

One prior art synchronization method involves the generation of a general system clock. Although a general system clock may be adequate for limited applications in the information handling system, it is not adequate for performance measurement purposes. For example, the system clock does not provide sufficient granularity for small (subcycle) variations between processors in the system. Further, the system clock requires a considerable amount of system resources which would adversely affect performance of the system if the system clock were continually referred to during performance measurement. This would further skew the performance results.

Some prior art approaches have attempted to solve the clock drift problem at the hardware level. For example, temperature data may be fed back to an oscillator, so that the oscillator can adjust its frequency to reduce or remove variations due to temperature.

Other prior art approaches have focused on clock synchronization in distributed systems, such as networks. Communications among computers connected by networks typically takes much longer than communications among tightly coupled processors in a multiprocessor system. Therefore, the accuracy requirement for clock synchronization in a distributed system is significantly less than the accuracy requirement for a multiprocessor system. The prior art synchronization methods used in distributed systems would not work with the required degree of accuracy for a multiprocessor system, especially if the multiprocessor system is a tightly couple, symmetrical, shared-memory multiprocessor system.

Consequently, it would be desirable to have a system and method of accurately and efficiently synchronizing and then mapping, or normalizing, processor clocks in a multiprocessor information handling system. It would be desirable to have a system and method which provide sufficient granularity for subcycle variations between processors, while taking into account the problem of clock drifts.

According to a first aspect of the present invention there is provided an information handling system, including: a plurality of processors, each processor having an internal clock; means for selecting a subset or all of the plurality of processors for the purposes of synchronization; means for synchronizing the internal clocks located on each selected processor; means for obtaining a time value to be mapped from a secondary clock located on a secondary processor; and means for mapping the time value to be mapped to an equivalent time value in a primary clock located on a primary processor.

The present invention is also directed to an information handling system capable of executing the above method, and to a computer readable medium including program code for implementing the above described method.

The present invention is then directed to a system and method for accurately and efficiently synchronizing and then mapping, or normalizing, processor clocks in a multiprocessor information handling system. The system and method of the present invention provide sufficient granularity for subcycle variations between processors, while taking into account the problem of clock drifts.

As will be described below in relation to the preferred embodiment, a plurality of processors in an information handling system is selected for the purpose of synchronization. The clocks located on the processors are synchronized, and then time values between synchronization points are mapped from each secondary processor to an equivalent, or normalized, time value in a primary processor. To accomplish this mapping, three clock differences are calculated. The first clock difference is the time between the first and second synchronization points for the primary processor, and the second clock difference is the time between the first and second synchronization points for the secondary processor. The third clock difference is the time difference between the time value to be mapped in the secondary processor and the first synchronization point in the secondary processor. The third time difference is multiplied by the ratio of the first time difference to the second time difference, and then the result is added to the time value of the primary clock at the first synchronization point. The final result is the mapped, or normalized, time value.

An advantage of the present invention is that any time value in any secondary processor may be mapped to an equivalent value in the primary processor. Another advantage of the present invention is that it substantially eliminates the problem of clock drifts in a multiprocessor information handling system. A further advantage of the present invention is that it provides the required degree of accuracy for use in a tightly coupled, symmetrical, shared-memory multiprocessor system.

A preferred embodiment of the invention will now be described, by way of example only, with reference to the accompanying drawings in which:

Figure 1 is a block diagram of an information handling system implementing the present invention;

Figures 2A and 2B are flow charts depicting a program executing on the system of Figure 1, and using the method of the present invention;

Figure 3 is a timing chart showing the synchronization of four processor clocks at three different synchronization points;

Figure 4 is a flowchart illustrating a method for mapping, or normalizing, processor clocks according to the present invention; and

Figure 5 is a timing chart showing the mapping, or normalization, of the processor clock readings from Figure 3.

Referring now to Figure 1, a multiprocessor information handling system in accordance with the present invention will be described. Information handling system 10 includes a plurality of processors 12. In the described embodiment, four processors 12 labeled P1, P2, P3, and P4 are included. Each of processors 12 is connected to memory 14 and I/O system 16 by system bus 18. Each of processors 12 may also be connected to cache 20, which may be either dedicated cache or shared cache, depending upon system requirements. The arrangement of memory system 14, I/O system 16, and cache 20 is not critical to the present invention and may be implemented in any number of well-known designs.

Each processor 12 includes an internal clock, represented as clock 22 in P1, clock 24 in P2, clock 26 in P3, and clock 28 in P4. Memory system 14 also includes one or more images of operating system 30 which controls the operation of information handling system 10. Operating system 30 may be any one of a number of well-known operating systems, such as the IBM OS/2 operating system.

Note that a multiprocessor information handling system with  $n$  processors, where  $n$  is the number of processors, is also referred to as an  $n$ -way multiprocessor system. Almost all of today's multiprocessor systems support "atomic" memory access operations. A memory operation is said to be "atomic" if the result of the operation is seen as a whole by any processor in the information handling system. As an example, if a shared memory variable, with an initial value of 100, is incremented by 10, the value of the shared variable is seen by any processor in the system as either 100 or 110, rather than as any other number. The present invention is preferably implemented on a system that supports atomic memory access operations, and in the following description, it is assumed that the system supports atomic memory access operations.

Note that processor clocks 22, 24, 26 and 28 may operate at slightly different and uneven speeds due to external factors, such as room temperatures. For example, suppose that the maximum clock speed variation of each processor is 0.5 ticks per million. On a multiprocessor system with processor clocks running at 100 MHz, during any one second period of time, any processor clock may be 50 ticks faster or slower than 100 MHz. In other words, if all processors start at the same clock value, during a one second period of time any two processors in the multiprocessor system might be 100 ticks apart at most.

A system and method for synchronizing plural processor clocks in a multiprocessor system is described in US patent application serial number 08/822,022. This application describes a method of reading processor clocks simultaneously in a multiprocessor information handling system. The basic synchronization method for synchronizing two processors requires that a shared variable, `sync_flag_1`, be set to zero before the process begins. After the shared variable `sync_flag_1` is set to zero, an iteration counter is set to a high number, such as 128 or 256. Next, another shared variable, `sync_flag_2`, is set to zero. `sync_flag_1` is incremented by one, and `sync_flag_1` is tested to determine if the value of `sync_flag_1` is equal to the number of processors being synchronized. If not, the process loops until `sync_flag_1` is equal to the number of processors being synchronized, and then the processor clock is read.

In this case, neither processor will proceed until both processors have incremented `sync_flag_1`. In other words, each processor increments `sync_flag_1` and then waits. Once each processor has incremented `sync_flag_1`, `sync_flag_1` will be equal to the number of processors being synchronized, i.e. in this example, `sync_flag_1` will be equal to two. This will cause both processors to read their respective internal clocks simultaneously, or nearly simul-

taneously. Note that the on-chip clock on each processor is also referred to as a cycle counter. The cycle counter typically runs at the speed of the processor. For example, in a 100 MHz processor, the cycle counter counts 100 million cycles, also referred to as ticks, per second. The number read from the internal clock is referred to as a cycle count, or as a number of ticks. After reading the processor clock, each processor stores the value read.

To increase the accuracy of the process, three separate sync flags are used, and the process is invoked three times. After the process is invoked three times, the iteration counter is decremented by one. A test is made to determine if the iteration count is zero. If the iteration count is zero, the average processor clock value is calculated, and synchronization of the two processor clocks is complete. If the iteration count is not equal to zero, the process continues until the iteration count reaches zero. The total number of times each processor reads its clock is the number of iterations multiplied by three (clocks are read three times per iteration). The resulting synchronized clock values are calculated by averaging the clock readings from multiple invocations of the clock synchronization process on each processor.

While the basic synchronization method may be used for synchronizing any number of processor clocks, the results become less accurate as more processor clocks are synchronized. To increase accuracy on systems with more than two processors, a more sophisticated method may be used as described below.

In those multiprocessor systems where there are more than two processors, one processor is designated as the primary processor and the rest are secondary processors. Note that any one of the processors may be chosen as the primary processor. Each secondary processor synchronizes with the primary processor by invoking the basic synchronization method as described above. Because the secondary processors synchronize with the primary processor at different times (i.e. one secondary processor at a time synchronizes with the primary processor), the results must next be adjusted. The purpose of the clock adjustment process is to relate all average processor clock values to a base average processor clock value, i.e. the average processor clock value of the primary processor.

The clock adjustment process for the primary processor first tests to determine whether the primary processor is to be synchronized with the first secondary processor. If yes, the process immediately exits because the primary processor and the first secondary processor have been previously synchronized. If the primary processor is to be synchronized with any secondary processor other than the first secondary processor, a clock adjustment is calculated. The clock adjustment is the difference between the average processor clock value of the primary processor when it is synchronized with the first secondary processor and the average processor clock value of the primary processor when it is synchronized with the current processor. The clock adjustment value is then written to a register for storing the clock adjustment value.

The clock adjustment process in each secondary processor first tests to determine if the selected secondary processor is the first secondary processor. If true, the clock adjustment process is exited for the reason stated above. If false, a test is made to determine if a clock adjustment value is ready. If not, the process loops until the clock adjustment value is ready. The clock adjustment value is ready when the primary processor has finished calculating it and has stored it in a known shared memory location, where it can then be read by the secondary processor. When the clock adjustment value is ready, the clock adjustment value is read, and the synchronized clock value of the selected secondary processor is adjusted.

As an example of applying the clock synchronization method described above, assume that clock readings of 600, 1200, 2000, and 1600 are returned for processors P1, P2, P3, and P4, respectively. This means that at some point in time, when the clock reading of processor P1 is 600, at exactly the same time the clock readings of processors P2, P3, and P4 are 1200, 2000, and 1600, respectively.

Referring now to Figure 2A, a program running in a multiprocessor information handling system that uses the present invention is described. The present invention is often used for tracing program execution in multiprocessor information handling systems. Each processor in the system starts execution at step 52 and exits at step 70. At step 52, a check is made to determine whether the program is finished. If it is, the program exits (step 70) on the particular processor. If the program is not finished, the processor performs some computation on behalf of the program (step 54). An event entry is logged and current time stamp is taken (step 56). The processor then returns to step 52 and the program continues to execute.

At any step within this execution loop, the processor can be interrupted, as shown in Figure 2B. For example, a system timer may interrupt each processor in the system at predetermined time intervals. When an interrupt occurs, a check is made (step 60) to determine whether a clock synchronization is needed by calculating the time difference between the current time and the last time clock synchronization was invoked. Typically, this check is performed by a system timer interrupt handler. If the system timer is running at 100 Hz, the system is interrupted by this timer once every 10 milliseconds. If clock synchronization is needed every second, then the clock synchronization should be invoked once every 100 times the system is interrupted by the system timer.

Assume that the maximum variation of clock speeds in a multiprocessor system is  $R_v$  ticks per second, and the time duration between two consecutive clock synchronizations is  $\Delta t$ . Thus the maximum clock variation between any two processors in the multiprocessor system during this period of time, represented as  $V_m$ , is described by the following equation:

$$V_m = 2 * R_v * D_t$$

Note that  $R_v$  is typically a constant if external conditions are constant. If  $V_m$ , the maximum clock variation between any two processors, is fixed, then the duration in which a clock synchronization is needed can be calculated by the following equation:

$$D_t = V_m / (2 * R_v)$$

There are certain times when  $R_v$  can not be assumed to be constant. For example, if a system is booting up from a cold state, then chip temperatures will rise rapidly, and the clock variation is relatively large. For situations where  $R_v$  is large, the clock synchronization has to be invoked more frequently, depending on the accuracy requirements of the application.

If  $D_t$ , the duration between two consecutive clock synchronization points, is fixed, then the maximum variation between any two processors can be calculated by the following equation:

$$V_m = D_t * 2 * R_v$$

For example, if clock speed varies by 5 ticks per second, then during a period of time of 10 seconds, the maximum clock difference between any two processor clocks is 100 ticks. On the other hand, if an application can tolerate 100 ticks of inaccuracy for the clocks, then the clock synchronization algorithm must be invoked at least once every 10 seconds.

Note that invocation of the clock synchronization method at interrupt time is one way of implementing the present invention. However, the method of the present invention may be implemented without using interrupts.

Still referring to Figure 2B, if a clock synchronization is required, the clock synchronization method of the present invention is invoked (step 62). This method is described in detail below, with reference to Figure 4. The synchronized time stamp is logged (step 64), and then appropriate interrupt handling is completed (step 66).

Referring now to Figure 3, assume that clocks P1, P2, P3, and P4 are synchronized at three points in time, as discussed above. At the first synchronization point 70, clocks P1, P2, P3, and P4 are synchronized at 600, 1200, 2000, and 1600, respectively. At the second synchronization point 72, they are synchronized at 700, 1310, 2120, and 1690, respectively. At the third synchronization point 74, they are synchronized at 800, 1430, 2330, and 1790, respectively.

Assume that the clock rates are constant between any two synchronization points. Therefore, clocks on processors other than P1 can be adjusted according to the rate on P1. In the example shown in Figure 3, once the clock readings have been synchronized, it is possible to normalize, or map, all clock readings on clocks P2, P3, and P4 to equivalent clock values on P1. The method for performing this adjustment, or mapping, is described below with reference to Figure 4.

Assume that a time value for a given processor clock P2 is to be mapped to a primary processor clock P1. Assume further that the time value to be mapped falls between two synchronization points, T0 and T1. The general method for mapping a clock value according to the present invention is as follows:

$$\text{Relative Value} = ((D_{t\_1}/D_{t\_2}) * (T - \text{Clock2\_T0})) + \text{Clock1\_T0}$$

$D_{t\_1}$  is the time difference between the two synchronization points, T0 and T1, in the primary processor clock (i.e. the processor clock to which the other processor clocks will be mapped).  $D_{t\_2}$  is the time difference between the same two synchronization points, T0 and T1, in the processor clock which is being mapped to the primary processor. T is the time value from P2 which is being mapped. Clock2\_T0 is the time value of clock P2 at time T0, and Clock1\_T0 is the time value of clock P1 at time T0.

Note that the error margin of the above mapping calculation is bounded by the maximum clock drift on processor P1 (i.e.  $R_v * D_{t\_1}$ ) between the two synchronization points T0 and T1. The error margin of the mapping calculation is defined as the difference between the mapped clock value on processor P1 and the actual equivalent clock value on P1. In other words, if two events on any two processors occur at two times, at least  $(2 * R_v * D_{t\_1})$  ticks apart, then the mapped clock values obtained from the above mapping algorithm will maintain the relative time order in which they occurred.

Referring now to Figure 4, a method for normalizing, or mapping, a clock reading will now be described. The first step is to determine if there is a clock reading, or time value, to be adjusted (step 80). If so, the time value is obtained

(step 82). Next, a test is performed to determine the source of the time value (i.e. is it a clock reading from P1, P2, P3, or P4?) (step 84). If the time value is from clock P1, no adjustment is needed, and the time value is output as is (step 86). If the time value is from a clock other than P1, an adjustment is necessary.

To adjust the time value to an equivalent P1 value first requires finding the synchronized period into which the time value falls (step 88). For example, a time value of 2060 on processor clock P3 falls in between the synchronized values of 2000 and 2120. The next step is to calculate the ratio of P1's difference to P3's difference (i.e.  $Dt\_1/Dt\_2$  in the equation above) (step 90). In the example given,  $Dt\_1$  is equal to 100 (i.e.  $700-600$ ),  $Dt\_2$  is equal to 120 (i.e.  $2120-2000$ ), and the ratio of  $Dt\_1$  to  $Dt\_2$  is thus  $100/120$ .

Next, the relative time value is calculated (step 92). In this example, T is equal to 2060, Clock1\_T0 is equal to 2000, and Clock0\_T0 is equal to 600. Thus, the relative time value is equal to  $((100/120) * (2060-2000)) + 600$ , which is equal to 650. Therefore, the adjusted time value, 650, is output in step 94.

Referring now to Figure 5, equivalent P1 time values are depicted for a time value of 2060 on processor clock P3, a time value of 1715 on processor clock P4, and a time value of 1370 on processor clock P2. The relative time values for P2 and P4 were obtained using the same method used for obtaining the relative time value for P3 (as described above with reference to Figure 4).

Once all four processors are synchronized, it is possible to calculate relative time values, which can be used for many purposes, including, but not limited to, performance measurement and analysis. As an example, suppose that a performance tool obtains the following system events regarding locks in the various processors in the information handling system. Locks are used extensively in multiprocessor systems to synchronize accesses to shared resources. Typically, a processor acquires a lock before using a system resource and releases the lock after finishing using it.

Processor	Time Stamp	Action
P1	785	Acquire
P1	790	Release
P2	1233	Acquire
P2	1255	Release
P2	1370	Acquire
P2	1394	Release
P3	2072	Acquire
P3	2096	Release
P4	1681	Acquire
P4	1701	Release

The first step is calculate an equivalent P1 time for each trace event. The equivalent times are calculated using the method described in Figure 4. In this example, equivalent times are as follows:

Processor	Time Stamp	Action	Equivalent P1 Time
P1	785	Acquire	785
P1	790	Release	790
P2	1233	Acquire	630
P2	1255	Release	650
P2	1370	Acquire	750
P2	1394	Release	770
P3	2072	Acquire	660
P3	2096	Release	680
P4	1681	Acquire	690
P4	1701	Release	710

Finally, the trace events can be ordered according to their equivalent P1 time, as depicted below. This ordering provides a picture of which events occurred, and in what order they occurred.

P1	P2	P3	P4	Equivalent P1 Time
	Acquire			630

(continued)

P1	P2	P3	P4	Equivalent P1 Time
Acquire Release	Release	Acquire Release	Acquire Release	650
				660
				680
				690
	Acquire Release			710
				750
				770
				785
				790

Although the present invention has been described with reference to performance tuning, the system and method of the present invention has many other uses. The invention may be used in any application where time stamps are logged and post processed. For example, operating systems may log the time of system events for trouble shooting purposes. Kernel level programs, such as the operating system kernel and device drivers, may use the present invention to assist with debugging time sensitive operations.

The system and method of the present invention is also not limited to mapping time between synchronization points. The invention may also be used to extrapolate out into the future. However, the maximum error margin an application can tolerate will determine how far out into the future the present invention can be used.

Although the invention has been described with a certain degree of particularity, it should be recognized that elements thereof may be altered by persons skilled in the art without departing from the spirit and scope of the invention. One of the preferred implementations of the invention is as sets of instructions resident in the memory 14 of one or more computer systems configured generally as described in Figure 1. Until required by the computer system, the set of instructions may be stored in another computer readable memory, for example in a hard disk drive, or in a removable memory such as an optical disk for eventual use in a CD-ROM drive or a floppy disk for eventual use in a floppy disk drive. Further, the set of instructions can be stored in the memory of another computer and transmitted over a local area network or a wide area network, such as the Internet, when desired by the user. One skilled in the art would appreciate that the physical storage of the sets of instructions physically changes the medium upon which it is stored electrically, magnetically, or chemically so that the medium carries computer readable information. The invention is limited only by the following claims.

## Claims

1. An information handling system, including:

a plurality of processors, each processor having an internal clock;

means for selecting a subset or all of the plurality of processors for the purposes of synchronization;

means for synchronizing the internal clocks located on each selected processor;

means for obtaining a time value to be mapped from a secondary clock located on a secondary processor; and

means for mapping the time value to be mapped to an equivalent time value in a primary clock located on a primary processor.

2. An information handling system as claimed in claim 1, wherein said means for mapping comprises:

means for determining a set of synchronization points;

means for determining one or more ratios between the synchronization points contained in the set; and

means for using the ratios to map the time value to be mapped to an equivalent time value in the primary clock

located on the primary processor.

3. An information handling system as claimed in claim 2, wherein said means for determining a set of synchronization points comprises means for determining a first synchronization point and a second synchronization point, wherein the time value to be mapped is between the first synchronization point and the second synchronization point.

4. An information handling system as claimed in claim 2, wherein said means for determining one or more ratios comprises:

means for calculating a first time difference, wherein the first time difference is a time between a first synchronization point in the set and a second synchronization point in the set for the primary processor;

means for calculating a second time difference, wherein the second time difference is a time between the first synchronization point in the set and the second synchronization point in the set for the secondary processor; and

means for calculating a ratio of the first time difference to the second time difference.

5. An information handling system as claimed in claim 2, wherein said means for using the ratios comprises:

means for multiplying the time value to be mapped by a ratio of a first time difference to a second time difference;

means for calculating a third time difference between the time value to be mapped and the time value of a secondary clock at a first synchronization point;

means for multiplying the third time difference by the ratio of the first time difference to the second time difference to obtain an intermediate result; and

means for adding a time value of the primary clock at the first synchronization point to the intermediate result to obtain the equivalent time value.

6. An information handling system as claimed in claim 2, further comprising:  
means for using the ratios to map the time value to be mapped to an equivalent time value in the primary clock located on the primary processor, wherein said time value to be mapped is a time value after the set of synchronization points.

7. An information handling system as claimed in any preceding claim, further including:

a memory system for storing instructions and data for use by the processors;

one or more images of an operating system for controlling the operation of the processors;

an input/output system for communicating information to and from peripheral devices; and

at least one system bus connecting the elements of the system for efficient operation.

8. A method for mapping a time value to be mapped from a secondary processor to an equivalent time value in a primary processor in a multiprocessor system, said method comprising the steps of:

selecting a plurality of processors for the purposes of synchronization;

synchronizing the internal clocks located on each selected processor;

obtaining a time value to be mapped from a secondary clock located on a secondary processor; and

mapping the time value to be mapped to an equivalent time value in a primary clock located on a primary processor.



9. A method for solving for clock drifts among a plurality of processors in a multiprocessor information handling system, comprising the steps of:

5 selecting a primary processor;

synchronizing a plurality of secondary processor clocks to a primary processor clock at a first synchronization point;

10 synchronizing the plurality of secondary processor clocks to the primary processor clock at a second synchronization point;

obtaining one or more time values to be mapped from the plurality of secondary clocks;

15 mapping the time values to be mapped to equivalent time values in the primary clock located on the primary processor; and

using the equivalent time values to correct for clock drift on a cycle by cycle basis in each secondary processor.

- 20 10. A method according to claim 9, wherein said mapping step comprises the steps of:

calculating a first time difference, wherein the first time difference is a time between a first synchronization point and a second synchronization point for the primary processor;

25 calculating a second time difference, wherein the second time difference is a time between the first synchronization point and the second synchronization point for the secondary processor;

calculating a ratio of the first time difference to the second time difference;

30 multiplying the time values to be mapped by a ratio of the first time difference to the second time difference;

calculating one or more third time differences between the time values to be mapped and the time value of the secondary clock at the first synchronization point;

35 multiplying each third time difference by the ratio of the first time difference to the second time difference to obtain one or more intermediate results; and

adding the time value of the primary clock at the first synchronization point to each intermediate result to obtain the equivalent time values.

- 40 11. A computer readable medium for mapping a time value to be mapped from a secondary processor to an equivalent time value in a primary processor in a multiprocessor system, comprising:

means for selecting a plurality of processors for the purposes of synchronization;

45 means for synchronizing the internal clocks located on each selected processor;

means for obtaining a time value to be mapped from a secondary clock located on a secondary processor; and

50 means for mapping the time value to be mapped to an equivalent time value in a primary clock located on a primary processor.




FIG. 1

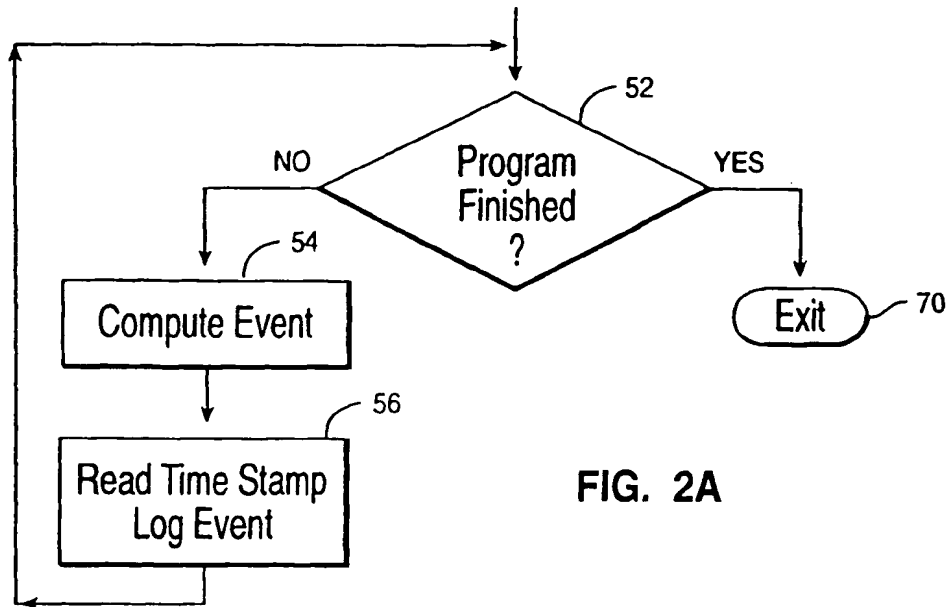


FIG. 2A




FIG. 2B




FIG. 3




FIG. 5

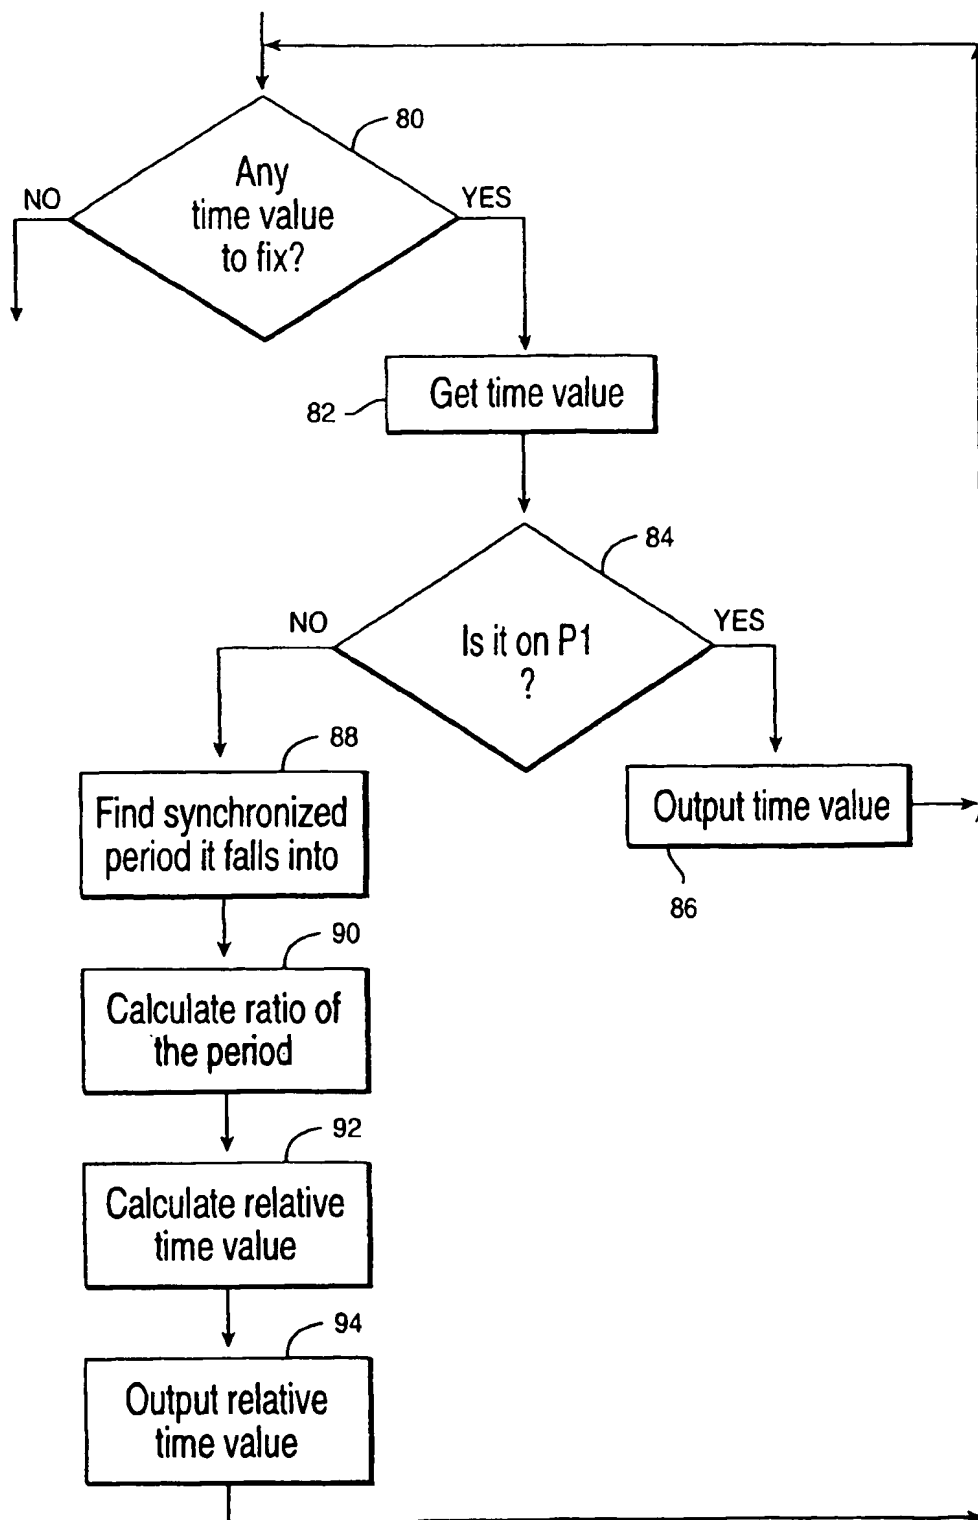


FIG. 4



European Patent  
Office

## EUROPEAN SEARCH REPORT

Application Number  
EP 98 30 4692

DOCUMENTS CONSIDERED TO BE RELEVANT			
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (Int.Cl.6)
A	KRISHNAN P ET AL: "A DISTRIBUTED REAL-TIME LANGUAGE AND ITS OPERATIONAL SEMANTICS" PROCEEDINGS OF THE REAL TIME SYSTEMS SYMPOSIUM, SANTA MONICA, DEC. 5 - 7, 1989, no. SYMP. 10, 5 December 1989, pages 41-50, XP000090987 INSTITUTE OF ELECTRICAL AND ELECTRONICS ENGINEERS * abstract * * page 47, left-hand column, last paragraph - right-hand column, line 34 * * figure 1 *	1,8,9,11	G06F1/14 G06F9/46
			TECHNICAL FIELDS SEARCHED (Int.Cl.6) G06F
The present search report has been drawn up for all claims			
Place of search THE HAGUE		Date of completion of the search 22 October 1998	Examiner Ciarelli, N
CATEGORY OF CITED DOCUMENTS X: particularly relevant if taken alone Y: particularly relevant if combined with another document of the same category A: technological background O: non-written disclosure P: intermediate document T: theory or principle underlying the invention E: earlier patent document, but published on, or after the filing date D: document cited in the application L: document cited for other reasons &: member of the same patent family, corresponding document			

EPO Form 1503 03/92 (P/4C01)